

Human Emulation: Progress Towards Realistic Synthetic Human Agents

Chris Forsythe

Patrick G. Xavier

Sandia National Laboratories*

Albuquerque, NM USA

jcforsy@sandia.gov

pgxavie@sandia.gov

ABSTRACT: *While there has been substantial progress in modeling and simulation of human agents, practical application is limited due to questions concerning the realism of agent behavior. At Sandia National Laboratories, a program of research and development is underway seeking a capability for highly realistic human emulation. Initially, emphasis was placed on computationally representing human Naturalistic Decision Making. The conceptual approach resulting from this work then served as a foundation for development of a framework for the comprehensive representation of decision processes. This framework utilizes a two-tiered architecture in which an underlying physiological model serves as the engine for a psychological model. Knowledge and cognitive processes are represented within the psychological model, whereas the physiological model provides the basis for incorporating organic factors (e.g., arousal, emotion, etc.). One output from the model is a simulated EEG signal. Exercises for validation have consisted of replicating studies that used human subjects and comparing actual results with simulated cognitive performance and EEG signals from the model.*

Ongoing work has two emphases. First, there is revision and extension of the model to accommodate practical applications. Second, knowledge representation capabilities are being expanded to incorporate massive volumes of knowledge and endow the model with human-like episodic memory. This latter development is especially important in that it creates the potential for customized agents that emulate specific cultures, or individuals. Current projects focus on modeling and simulation of human cognition and behavior. Specific applications include small unit combat, Insider threats and human performance in high consequence systems. This work also emphasizes the utilization of human cognitive models within the context of intelligent systems.

Keywords:

Cognitive modeling, naturalistic decision making, recognition-primed decision making.

1. Introduction

For some time, synthetic humans and other synthetic organic entities have been an important component for simulation and gaming systems. Varying levels of sophistication may be observed in the underlying cognitive models. For instance, by increasing the volume and breadth of procedural knowledge, larger ranges of behavioral response may be attained. Likewise, through incorporation of computational approaches such as AI-based reasoning, learning, fuzzy logic, neural nets and

genetic algorithms, increased flexibility and adaptability may be achieved.

It is important to note that a broad range of behavioral response and a high level of flexibility and adaptability do not equate to realism. However, with the exception of cinematic computer graphics, most current applications impose minimal demand for behavioral realism in synthetic entities. For instance, entertainment is the primary consideration with most gaming applications. Thus, there has been emphasis on visual, audio and other features that enhance the user experience. Similarly, in

* This work was performed at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

current military simulation systems, attention has often been focused on replicating characteristics of equipment so as to provide a high fidelity simulation of actual operational experience. As a result, an emphasis has been placed on creating a sensory-motor experience.

Many simulation applications introduce substantially greater demands for behavioral realism in synthetic entities than has been typical. For instance, simulation used for assessment of high consequence applications (e.g., security, tactical effectiveness, operator response, etc.) requires confidence that synthetic entities are behaving in a manner comparable to actual humans placed in similar circumstances. For training military, law enforcement and other personnel in situations that are highly ambiguous and involve personal interactions, synthetic entities should exhibit behavioral variability consistent with actual populations. Furthermore, to prepare for specific confrontations (e.g., individual, culture, group), there is need for synthetic entities endowed with a rich personal history that may be customized to fit the situation.

This paper summarizes a program of research and development undertaken by Sandia National Laboratories for the development of a cognitive architecture that will enable highly realistic, customizable synthetic entities. This paper discusses three facets of this work: (1) the computational representation of human naturalistic decision making; (2) incorporation of organic factors (e.g., stress, fatigue, emotion, etc.); and (3) endowment of synthetic entities with human-like episodic memory.

2. Reverse Engineering Human Cognition

2.1 General Approach

In developing a cognitive architecture, an approach has been taken that may be characterized as reverse engineering. While at a descriptive level there was a desire to support Human Naturalistic Decision Making, initially, certain fundamentals were accepted. For example, a heavy emphasis has been placed on oscillating systems theory as an explanation for the mechanisms underlying semantic and episodic memory (Klimesch, 1996). Given these fundamentals, published research with human subjects has provided the basis for creating a set of design specifications. These specifications tend to be of the form, "if input x is applied, output y should be observed." In generating specifications, attention has been focused on capturing the relationships between cognitive performance and electrophysiological phenomena. One output of the cognitive architecture is a simulated electroencephalograph signal. Typical test

conditions present the emulator a range of stimulus conditions with the objective being the design of a cognitive architecture that behaves in accordance with observations from human subjects. For example, Niedeggen & Rosler (1999) reported increased amplitude of response in event-related potentials relative to the spreading activation generated by stimulus concepts. In testing the simulation, concepts producing low medium and high levels of spreading activation were presented and the expected difference in response amplitude demonstrated.

2.2 Simulation of Human Naturalistic Decision Making

Klein and others (Zsombok & Klein, 1997) have provided descriptive models of the process by which expert decision makers arrive at decisions in realistic settings. Currently, the simulation provides a computational representation of Level 1 decision making from Klein's Recognition-Primed Decision (RPD) making model (Klein, 1997). Here, within ongoing events, the decision maker recognizes a pattern of cues associated with a known "situation." Once this recognition occurs, there is implicit knowledge of the appropriate course of action, as well as goals and expectations.

Figure 1 provides a conceptual depiction of how RPD has been represented computationally. The synthetic entity is attributed knowledge of situations. For the illustrated application, the emphasis was Close Quarters Battle, and situations consisted of tactics for exterior movement. Environmental cues and related knowledge create patterns in an associative semantic activation network, and when these patterns corresponded to the pattern template associated with a known tactic, the situation is recognized as one appropriate for that tactic. In the prototyped application, heuristics then provide agents generic instructions for implementing the tactic.

2.3 Incorporation of Organic Factors

While the initial model depicted in Figure 1 offered a computational representation based on RPD, the resulting decision maker was an idealized, "perfect decision maker." By this, it is meant that the behaviors of the decision maker and its decision making processes were strictly determined by its inputs and knowledge, instead of being affected by fear, arousal, stress, etc. Furthermore, agents exhibited no individual differences attributable to cultural factors or personal experiences. Collectively, these factors have been termed "organic factors" and a model has been advanced to account for their influence on systems (Forsythe & Wenner, 2000).

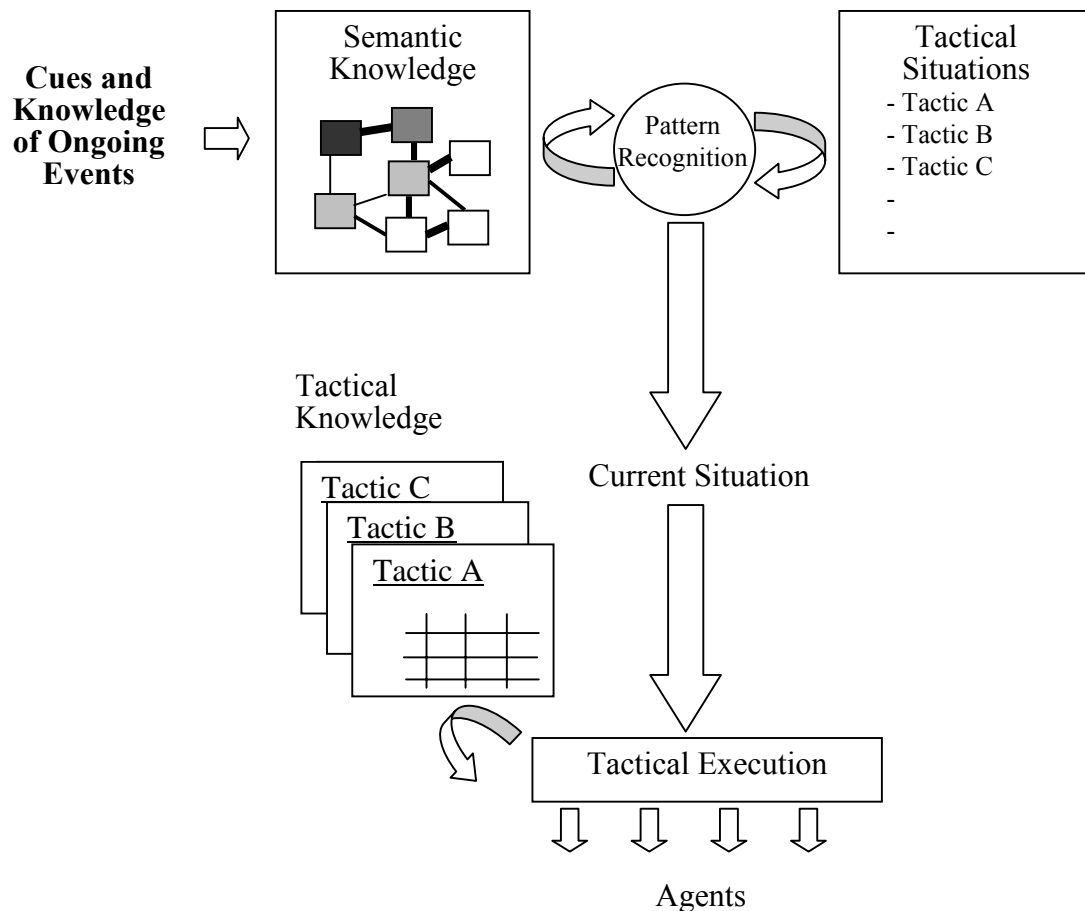


Figure 1. Conceptual Framework for implementation of Level 1 Recognition Primed Decision Making in a Computer Simulation.

Early in the model's development, it was realized that a purely psychological model, exemplified by Figure 1, would be inadequate for representing the influence of organic factors on cognitive behavior. There is enormous ambiguity in basic terminology (e.g., stress, arousal) and without a representation of underlying mechanisms, the scope and predictive capabilities would be severely limited. However, many facets of cognitive behavior (e.g., knowledge representation) are well described by psychological models (Goldsmith et.al., 1991). Consequently, a two-tiered approach was adopted in which knowledge is represented using a psychological model, while a separate physiology-based model serves as the engine that drives this psychological model. (See Figure 2.) The fact that knowledge is not directly represented in the neural (i.e., physiological) model distinguishes this design from neural net and connectionist approaches, yet facilitates representation of

the vast quantities of knowledge essential to a realistic emulation.

The mapping of the psychological to the physiological model was critical. Concepts embodied by our earlier instantiation of RPD were retained. This included a separate representation of individual situational elements, pattern recognition and activation of schema-like representation of known situations. Frame/Content theory provided an initial bridge. This theory asserts that the representation of individual elements of content within a structural or contextual frame is a basic organizing principle of the neural system (MacNeilage, 1998). Examples include figure/ground relationships in perception, syntax and semantics in linguistics, and differential motor specialization for stabilization and manipulation. Generalizing frame/content theory,

individual elements of a situation represent content, whereas situation schema provide an interpretive frame.

Further extension involved mapping these ideas to the model of memory processes proposed by Wolfgang Klimesch and colleagues (Klimesch, 1996). Two phenomena have been described. First, in the absence of intrinsic or extrinsic stimulation, regions associated with semantic memory exhibit synchronous activation in the high alpha (10-13 Hz) bandwidth. Once stimulated, desynchronization occurs. It is suggested that semantic memory processes involve the activation of numerous localized neural assemblies. (Neural assemblies contain a collection of individual neural units with the operation of individual units dictated by low-level neural processes, e.g., transmitter-receptor interactions, metabolic properties, etc.) These assemblies oscillate in phase with pulses from a pacemaker until stimulated, at which time activation increases and assemblies begin to oscillate independent of the pacemaker. At this point, there is desynchronization. In contrast, episodic processes exhibit a completely different profile. Specifically, processing demands lead to increased synchronization in the theta (4-7 Hz) bandwidth. This pattern of activation is consistent with oscillation of a single distributed neural assembly.

These ideas are crucial to the current mapping of psychological to physiological processes. In particular, activation associated with individual elements of a situation is equated to the activation of numerous localized assemblies with oscillations in the 10-13 Hz bandwidth. Simultaneously, there is a separate pattern recognition process that monitors activation of assemblies associated with individual elements and responds when specified patterns of activation occur. This would be synonymous with matching current conditions to a known situation schema. The pattern recognition process is associated with a single neural assembly that oscillates in the 4-7 Hz bandwidth.

As illustrated in Figure 2, a semantic activation network is used to represent semantic knowledge activated by individual elements of a situation. This network consists of nodes for individual concepts, and associative links between nodes that differ in their strength of association. Each concept in the psychological model is assigned a separate neural assembly, and the activation of each concept is a function of the activation of the neural assembly assigned to it. In the computational model, the update frequency of each individual component (e.g., a concept node, situation recognizer, etc.) is the theoretical oscillation frequency of its physiological analogue.

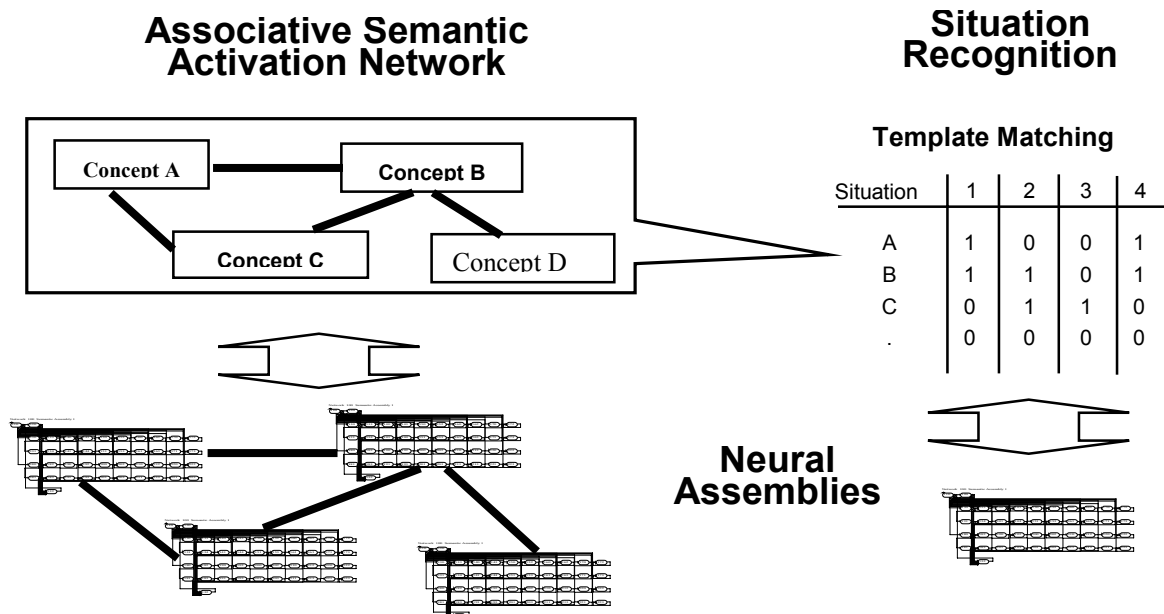


Figure 2. Framework for Emulator showing psychological and physiological levels.

Situation recognition was initially represented in the psychological model by a template matching process. Rows of the template represent known situation schema and columns correspond to concepts in the knowledge network. A simplified approach may be utilized whereby binary numbers indicate the activation, or lack of activation, of individual concepts during a given time period. Recognition occurs incrementally in accordance with a race model and when a threshold is exceeded, there is activation of the situation schema. This approach, however, is considered overly simplistic and current efforts are focused on better delineating the process by which humans recognize situations and computationally representing those processes (e.g., evidence accumulation and neural net approaches are both being explored).

With this model, employing the reverse engineering process described earlier, organic factors may be incorporated into the simulation. For example, arousal may be manipulated by adjusting the pulse rate of the pacemaker. In tests wherein arousal was manipulated in this manner, the anticipated effects on electrophysiological response and cognitive performance were observed. Emotion has been addressed through instantiation of the model proposed by Joseph LeDoux (1998). Here, concepts and situations may be assigned emotional associations. Their activation leads to activation of emotional processes. The result is heightened activation of the initial concept, and associated concepts, and active inhibition of unrelated concepts. (See Figure 3). As illustrated by the examples of arousal and emotion, the model offers hooks that also allow for the effects of other organic factors (e.g., fatigue, metabolic, psychotropic substances, etc.) to be introduced.

3. Simulation of Episodic Memory

Generally, the knowledge endowed to synthetic entities has been restricted to that directly relevant to the application domain. For instance, a synthetic fighter pilot knows about air combat and nothing else. In reality, individuals have a collective life experience that may exert as strong, and sometimes stronger, influence on decisions than domain knowledge. Furthermore, differences in personal experience are a primary factor accounting for the individual variability observable in human behavior. For example, individuals may interpret similar events differently on the basis of the contexts each individual has previously experienced a similar event (e.g., exposure to an ethnic group through television versus real life). For these reasons, it is believed that a vital element in creating realistic synthetic entities will involve the ability to endow agents with a synthetic life history.

Current efforts have focused on developing a capability that allows synthetic entities to meaningfully represent their experiences. A conceptual intelligent machine example illustrates the direction of this work. Here, two robotic vehicles systematically search a building to locate a smoke source. Based on their sensors, communications and data processing capabilities, as they progress through the scenario, different concepts in their semantic networks are activated (See Figure 4a). The result is a time series of patterns of semantic activation. This time series may be statistically analyzed to identify recurrent schema (e.g., progressing down a hallway following a smoke gradient). This is illustrated in Figure 4b. Endowed with knowledge of these schema, stories may be constructed that are based on the sequence of schema experienced during a given event (See Figure 4c). Given knowledge of these schema, subsequent analysis allow identification of recurrent sequences of schema (i.e., themes or storylines).

Preliminary results on this branch of our work are presented in (Schoenwald, et. al. 2002), which examines the use of computer simulation of a collective of embodied agents to generate episodic memory. In the simulated scenario, ground-based robotic vehicles attempt to search a building for smoke and place a robot at the highest smoke concentration the collective finds. The simulations, done in Sandia's Umbra embodied agent simulation environment (Gottlieb, et. al., 2001, 2002), included the behavior, sensing, control, mobility, and communication of a seven robot team. (See Figure 5.)

As in the conceptual description, the robots' logical and sensor states activate nodes in semantic network (30 concepts) whose outputs were monitored to generate trace data, in this case for 20 five-minute runs. Two cluster analysis methods, K-means and DIANA, were applied to randomly sampled activation data for a 15-concept subset chosen by a domain non-expert. Although the concepts selected did not include several concepts domain experts would have chosen, the K-means analysis produced clusters consistent with high-level themes (roles) identified by one of the robot programmers. Continuation of this work will include improvements to activation tracing and selection of concepts, as well as using the clusters to enable the robots to express their activity at the role level.

The capabilities described here are an initial step toward endowing synthetic entities with a life history. For instance, parallel efforts focus on mechanisms for generating unique life histories. It should be noted that these capabilities provide the basis for mental simulation, a key ingredient for expanding the model of RPD to levels 2 and 3 (Klein, 1997). Furthermore, the episodic

representations also provide a basis for synthetic entities to learn over the course of simulator runs and even develop knowledge based on shared experiences with individual trainees. Consequently, a synthetic entity might remind a trainee what happened in a similar scenario two months earlier or exhibit differential confidence based on the trainees' recent success or failure.

4. Toward Practical Application

Sections 1 & 2 described a psycho-physiologically motivated implementation of the Level 1 RPD model applied to decision-making disembodied from the world, with perception and action abstracted away. We now consider a direction for integrating our implementation of the RPD core into the behavior of embodied agents, particularly simulated Computer Generated Forces and (real) autonomous robotic systems.

Recall that RPD is a descriptive model of how people determine or update a Course of Action (COA) in the context of dynamic situations. This leads to the question of *what* a COA is in a computational sense. (I.e., *is it a script? A planning mechanism?*) In an embodied agent, a COA is simply some mechanism that produces action. In our view, it is important that such a mechanism may take into account the agent's perception of and model (e.g., cognitive map) of the world, its perception of its own state, its cognitive state, and that mechanism's own internal state. Furthermore, in the embodied agent context, *action* must be interpreted generally, so that an action can control sensors, alter perceptual processing, update cognitive state, or attempt to cause changes to physical state through motor actions. Figure 6 shows a way in which the RPD implementation of the previous sections might be integrated into an embodied agent.

In the context of our particular implementation of Situation Recognition, Perception (in the general sense, which can also be applied to a world model) may cause node (concept) stimulation in the semantic activation network and read node activations. Thus, Perception can activate cues for situations, and semantic activation may play a role in "top-down" influences on Perception. COA mechanisms not only may cause and observe semantic activation, but they may inhibit activation. Thus, if state machines (finite, fuzzy finite, hierarchical, etc.) are used as COA mechanisms, the states of their finite control may be integrated with semantic activation, or at least be reflected by it. (COA mechanisms may also be based on other approaches, such as ACT-R or SOAR.) We leave open the possibility that situations may be nested and that a COA might entail concurrent behaviors or mechanisms.

Resolving conflicts that might arise can be viewed as a part of Level 2 and 3 RPD; alternatively, conflicts might be viewed as situations themselves, and mechanisms for resolving them, such as applications of rule-based cognition, may be viewed as part of an appropriate COA.

5. Conclusion

This paper has described a theoretical approach and concepts for the development of realistic synthetic humans. A limited capability for modeling somewhat constrained scenarios is currently available and work now focuses on expanding this capability for modeling more open-ended dynamic scenarios. It is acknowledged that the level of realism attained may be debated, however it is believed the approaches described here offer meaningful progress relative to current alternatives. In closing, it is noted that putting aside the debate concerning realism and synthetic humans, these same capabilities may be applied to intelligent machine applications with equal and perhaps more immediate advantages.

6. References

- Forsythe, C. & Wenner, C.A. (2000). Surety of human elements of high consequence systems: An organic model. In *Proceedings of the IEA 2000/ HFES 2000 Congress*, 3-839 – 3-842.
- Goldsmith, T.E., Johnson, P.J. & Acton, W.H. (1991). Assessing structural knowledge. *Journal of Educational Psychology*, 83(1), 88-96.
- Gottlieb, E., McDonald, M. & Oppel, F. (2002) Umbra's High Level Architecture (HLA) Interface. SAND 2002-XXXX, Sandia National Laboratories.
- Klein, G. (1997). An overview of naturalistic decision making applications. In C.E. Zsombok & G. Klein *Naturalistic Decision Making*, Mahwah, NJ: Lawrence Erlbaum, 49-59.
- Klimesch, W. (1996). Memory processes, brain oscillations and EEG synchronization. *International Journal of Psychophysiology*, 24, 61-100.
- LeDoux, J. (1998). The emotional brain: The mysterious underpinnings of emotional life. Touchstone Books.
- MacNeilage, P.F. (1998). The frame/content theory of evolution of speech production. *Behavioral and Brain Sciences*, 21, 499-546.

Niedeggen, M. & Rosler, F. (1999). N400 effects reflect activation spread during retrieval of arithmetic facts. *Psychological Science*, 10(3), 271-276.

Schoenwald, D., Xavier, P., Thomas, E., Forsythe, C., & Parker, E. (2002). Simulation of a Cognitive Algorithm for a Distributed Robotic Sensing Network. In *Proceedings of the 9th International Symposium on Robotics and Applications*, Orlando, FL, June 2002.


Zsombok, C.E. & Klein, G. (1997). *Naturalistic Decision Making*, Mowah, NJ: Lawrence Earlbaum.

Author Biographies


J. CHRIS FORSYTHE, PhD, is a Principal Member of Technical Staff in the Computational Initiatives

organization at Sandia National Laboratories. Current projects focus on modeling and simulation of human cognition and behavior. Specific applications include small unit combat, Insider threats and human performance in high consequence systems. This work also emphasizes the utilization of human cognitive models within the context of intelligent systems (i.e., cognitive systems).

PATRICK G. XAVIER, PhD, is a Principal Member of Technical Staff in the Intelligent Systems Principles organization at Sandia National Laboratories. Ongoing activities include the simulation of heterogeneous systems including embodied agents and mobile robots, including sensing, behavior, cognition, communications, and motion control. Recent work has also included automated and interactive motion-planning, as well as distance computation and collision-detection for simulation and motion planning.




Concepts Not
Associated with Fear
Inducing Stimulus




Concepts Associated
with Fear Inducing
Stimulus

Activation in Context without Fear Inducing Stimulus



Concepts Not
Associated with Fear
Inducing Stimulus



Concepts Associated
with Fear Inducing
Stimulus

Activation in Context with Fear Inducing Stimulus

Figure 3. Differential Response to Conditions with and without Fear Inducing Stimulus




Figure 4a. Example Time Series of Patterns of Semantic Activation

	t ₁	t ₂	t ₃	t ₄	t ₅	t ₆	t ₇	t ₈	t ₉	t ₁₀	t ₁₁	t ₁₂	t ₁₃	t ₁₄	t ₁₅	t ₁₆
R1 Smoke	0	0	0	0	1	2	3	4	4	3	3	3	4	4	4	5
R1 Passage	0	1	0	0	0	0	0	1	0	0	0	0	0	0	1	0
R1 Hallway	0	0	1	1	1	1	1	0	1	1	1	1	0	0	0	1
R1 Intersection	0	0	0	0	0	0	1	0	0	0	0	0	1	1	0	0
R2 Alarm	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
R2 Smoke	0	0	0	0	0	1	3	5	7	8	8	8	8	7	8	8
R2 Passage	0	1	1	0	0	0	0	0	0	0	0	1	0	0	0	0
R2 Hallway	0	0	0	1	1	1	1	1	0	0	0	0	1	1	1	0
R2 Intersection	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	1
R1-R2 Direction	1	1	0	1	1	1	1	1	1	1	0	0	0	1	0	0
R1-R2 Separation	1	1	5	5	5	5	5	5	5	5	5	5	5	5	5	4

Figure 4b. Example Derivation of Schema Based on Recurrent Patterns of Semantic Activation

1. Entered building
2. Searched for smoke, found no smoke
3. Selected path, passage into hallway
4. Followed path (search smoke)
5. Detected smoke
6. Followed path (smoke gradient), reached intersection
7. Sampled paths, found path with more smoke
8. Followed path (smoke gradient), reached intersection
9. Alerted (destination)
10. Followed path (destination)

Figure 4c. Story Generated Based on Sequential Ordering of Schema at the Conclusion of Simulation Run.




Figure 5. Detailed simulation (Umbra) of multiple vehicles navigating a building without maps or localization. The robot behaviors and sensor suites were modified to place a robot at the maximum smoke level encountered. The left view shows a cut-away view of the building. The right view shows a close-up of vehicles with their IR sensing visible.

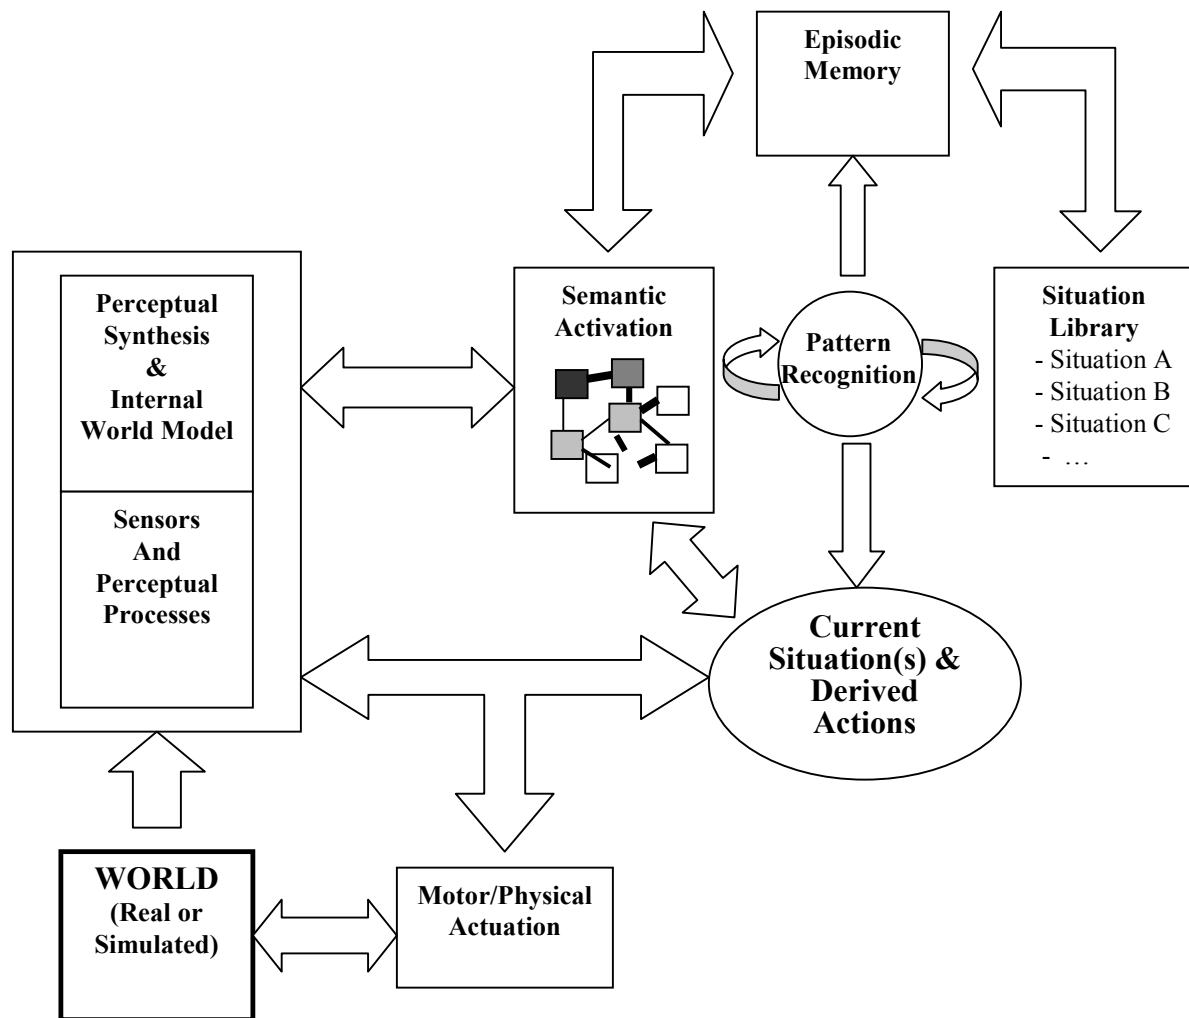


Figure 6. A simplified diagram showing conceptually how an RPD model can be integrated into a framework for embodied agent behavior.